

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	7951	"storage area network"	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:42
S2	10	barroux.in.	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:44
S3	176	libert.in.	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:46
S4	186	S2 S3	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:43
S5	5161	content with distribution with network	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:44
S6	54610	"709"/\$.ccls.	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:44
S7	1043	S5 and S6	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:44
S8	13614387	@ad<"20030630"	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:44
S9	662	S7 and S8	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:44
S10	1213902	access	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:44
S11	613	S9 and S10	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:45

EAST Search History

S12	289	router and S11	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:45
S13	213	traffic and S12	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:45
S14	168	bandwidth and S13	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:46
S15	8222	(preferred optimal) with path	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:46
S16	12716	(preferred optimal) with (path route)	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:46
S17	25	S16 and S9	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/05/31 20:47



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: ☒ The ACM Digital Library ☐ The Guide

+"content distribution network" proximity qos priority "routing



THE ACM DIGITAL LIBRARY



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Published before June 2003

Terms used content distribution

Found 38 of 142,723

network proximity qos priority routing path optimal weighted

Sort results
by

relevance



[Save results to a Binder](#)

Try an [Advanced Search](#)

Display
results

expanded form



[Search Tips](#)

Try this search in [The ACM Guide](#)

☐ Open results in a new
window

Results 1 - 20 of 38

Result page: [1](#) [2](#) [next](#)

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [OverQoS: offering Internet QoS using overlays](#)



Lakshminarayanan Subramanian, Ion Stoica, Hari Balakrishnan, Randy H. Katz

January 2003 **ACM SIGCOMM Computer Communication Review**, Volume 33 Issue 1

Publisher: ACM Press

Full text available: [pdf\(722.77 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper proposes *OverQoS*, an architecture for providing Internet QoS using overlay networks. OverQoS empowers third-party providers to offer enhanced network services to their customers using the notion of a *controlled loss virtual link* (CLVL). The CLVL abstraction bounds the loss-rate experienced by the overlay traffic; OverQoS uses it to provide differential rate allocations, statistical bandwidth and loss assurances, and enables explicit-rate congestion control algorithms.

2 [Consistency and replication: Modeling redirection in geographically diverse server sets](#)



Lisa Amini, Anees Shaikh, Henning Schulzrinne

May 2003 **Proceedings of the 12th international conference on World Wide Web WWW '03**

Publisher: ACM Press

Full text available: [pdf\(362.44 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Internet server selection mechanisms attempt to optimize, subject to a variety of constraints, the distribution of client requests to a geographically and topologically diverse pool of servers. Research on server selection has thus far focused primarily on techniques for choosing a server from a group administered by single entity, like a content distribution network provider. In a federated, multi-provider computing system, however, selection must occur over distributed server sets deployed by ...

Keywords: content distribution network (CDN), performance, server selection, web traffic redirection

3 [Network behavior: The effectiveness of request redirection on CDN robustness](#)

Limin Wang, Vivek Pai, Larry Peterson

 December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Publisher: ACM Press

Full text available:  pdf(1.86 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

It is becoming increasingly common to construct network services using redundant resources geographically distributed across the Internet. Content Distribution Networks are a prime example. Such systems distribute client requests to an appropriate server based on a variety of factors---e.g., server load, network proximity, cache locality--in an effort to reduce response time and increase the system capacity under load. This paper explores the design space of strategies employed to redirect reque ...

4 High-density model for server allocation and placement



 Craig W. Cameron, Steven H. Low, David X. Wei

June 2002 **ACM SIGMETRICS Performance Evaluation Review**, **Proceedings of the 2002 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '02**, Volume 30 Issue 1

Publisher: ACM Press

Full text available:  pdf(107.29 KB)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

It is well known that optimal server placement is NP-hard. We present an approximate model for the case when both clients and servers are dense, and propose a simple server allocation and placement algorithm based on high-rate vector quantization theory. The key idea is to regard the location of a request as a random variable with probability density that is proportional to the demand at that location, and the problem of server placement as source coding, i.e., to optimally map a source value (r ...

Keywords: content distribution, high density, server placement and allocation

5 Mobility and Wireless Access: Mobile streaming media CDN enabled by dynamic



 **SMIL**

Takeshi Yoshimura, Yoshifumi Yonemoto, Tomoyuki Ohya, Minoru Etoh, Susie Wee
May 2002 **Proceedings of the 11th international conference on World Wide Web WWW '02**

Publisher: ACM Press

Full text available:  pdf(623.98 KB)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper, we present a mobile streaming media CDN (Content Delivery Network) architecture in which content segmentation, request routing, pre-fetch scheduling, and session handoff are controlled by SMIL (Synchronized Multimedia Integrated Language) modification. In this architecture, mobile clients simply follow modified SMIL files downloaded from a streaming portal server; these modifications enable multimedia content to be delivered to the mobile clients from the best surrogates in the CD ...

Keywords: CDN, SMIL, mobile network, streaming media


6 Session 1: DNS: King: estimating latency between arbitrary internet end hosts



 Krishna P. Gummadi, Stefan Saroiu, Steven D. Gribble

November 2002 **Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement IMW '02**

Publisher: ACM Press

Full text available:  pdf(1.40 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The ability to estimate network latencies between arbitrary Internet end hosts would enable new measurement studies and applications, such as investigating routing path inefficiencies on a wide-scale or constructing topologically sensitive overlay networks: In this paper we present King, a tool that accurately and quickly estimates the latency between arbitrary end hosts by using recursive DNS queries in a novel way. Compared to previous approaches, King has several advantages. Unlike IDMaps, Ki ...

Keywords: latency measurement tool, recursive DNS

7 Characterizing the scalability of a large web-based shopping system



August 2001 **ACM Transactions on Internet Technology (TOIT)**, Volume 1 Issue 1

Publisher: ACM Press

Full text available: pdf(261.91 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This article presents an analysis of five days of workload data from a large Web-based shopping system. The multitier environment of this Web-based shopping system includes Web servers, application servers, database servers, and an assortment of load-balancing and firewall appliances. We characterize user requests and sessions and determine their impact on system performance scalability. The purpose of our study is to assess scalability and support capacity planning exercises for the multit ...

Keywords: capacity planning, clustering, personalization, scalability analysis, web-based systems, workload characterization

8 Topology and routing: Topology modeling via cluster graphs



Balachander Krishnamurthy, Jia Wang

November 2001 **Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement IMW '01**

Publisher: ACM Press

Full text available: pdf(782.02 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Several recent studies have focused on generating Internet topology graphs. Topology graphs have been used to predict growth patterns of prefixes and traffic flow as well as for designing better protocols. Internet topology graphs can be studied at either inter-domain level or router level. For some applications, inter-domain level topology graph is too coarse, while router level topology graph may be too fine-grained. We introduce *cluster graphs* as a way of modeling Internet topology at ...

9 Energy-conserving data placement and asynchronous multicast in wireless sensor networks



Sagnik Bhattacharya, Hyung Kim, Shashi Prabh, Tarek Abdelzaher

May 2003 **Proceedings of the 1st international conference on Mobile systems, applications and services MobiSys '03**

Publisher: ACM Press

Full text available: pdf(222.77 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

In recent years, large distributed sensor networks have emerged as a new fast-growing application domain for wireless computing. In this paper, we present a distributed application-layer service for data placement and asynchronous multicast whose purpose is power conservation. Since the dominant traffic in a sensor network is that of data retrieval, (i) caching mutable data at locations that minimize the sum of request and update traffic, and (ii) asynchronously multicasting updates from sensors ...

10 Dynamic services and analysis: Engineering and hosting adaptive freshness-sensitive web applications on data centers



Wen-Syan Li, Oliver Po, Wang-Pin Hsiung, K. Selçuk Candan, Divyakant Agrawal
May 2003 **Proceedings of the 12th international conference on World Wide Web WWW '03**

Publisher: ACM Press

Full text available: pdf(10.31 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Wide-area database replication technologies and the availability of content delivery networks allow Web applications to be hosted and served from powerful data centers. This form of application support requires a complete Web application suite to be distributed along with the database replicas. A major advantage of this approach is that dynamic content is served from locations closer to users, leading into reduced network latency and fast response times. However, this is achieved at the expense ...

Keywords: database-driven web applications, dynamic content, freshness, response time, net-work latency, web acceleration

11 Performance Workload Char. and Adaptation: Improving web performance by client characterization driven server adaptation



Balachander Krishnamurthy, Craig E. Wills
May 2002 **Proceedings of the 11th international conference on World Wide Web WWW '02**

Publisher: ACM Press

Full text available: pdf(241.76 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We categorize the set of clients communicating with a server on the Web based on information that can be determined by the server. The Web server uses the information to direct tailored actions. Users with poor connectivity may choose not to stay at a Web site if it takes a long time to receive a page, even if the Web server at the site is not the bottleneck. Retaining such clients may be of interest to a Web site. Better connected clients can receive enhanced representations of Web pages, such ...

Keywords: client characterization, client connectivity, server adaptation

12 Scalability study of the ad hoc on-demand distance vector routing protocol



Sung-Ju Lee, Elizabeth M. Belding-Royer, Charles E. Perkins
March 2003 **International Journal of Network Management**, Volume 13 Issue 2

Publisher: John Wiley & Sons, Inc.

Full text available: pdf(669.50 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

As mobile networking continues to experience increasing popularity, the need to connect large numbers of wireless devices will become more prevalent. Many recent proposals for ad hoc routing have certain characteristics that may limit their scalability to large networks. This paper examines five different combinations of modifications that may be incorporated into virtually any on-demand protocol in order to improve its scalability. The scalability of current on-demand routing protocols is evalu ...

13 Session 5: WWG: a wide-area infrastructure to support groups



Joan Manuel Marquès, Leandro Navarro
September 2001 **Proceedings of the 2001 International ACM SIGGROUP Conference on**

Supporting Group Work GROUP '01



Publisher: ACM Press

Full text available: pdf(535.53 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Group learning at Internet scale is becoming more frequent in university courses. This complex process requires support by distributed computing learning support infrastructures. This paper describes the design of WWG (World-Wide Groups): a distributed and decentralized infrastructure with the aim of supporting distributed group learning and team work, centered on the distribution of events, so that every participant can be notified and thus be aware of the actions, changes, progress of the group ...

Keywords: CSCL environment, distance cooperative learning, event distribution, virtual groups

14 Extended abstracts: Operating system support for massive replication



Arun Venkataramani, Ravi Kokku, Mike Dahlin

July 2002 **Proceedings of the 10th workshop on ACM SIGOPS European workshop: beyond the PC EW10**

Publisher: ACM Press

Full text available: pdf(293.50 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

The increasing number of devices used by each user to access data and services and the increasing importance of the data and services available electronically both favor "access-anywhere" network-delivered services. Unfortunately, making such services highly available is difficult. For example, even though end servers or service hosting sites advertise an availability of "four nines" (99.99%) or "five nines" (99.999%), the end-to-end service availability (as perceived by clients) is typically limited ...

15 Application level performance: On the use and performance of content distribution networks



Balachander Krishnamurthy, Craig Wills, Yin Zhang

November 2001 **Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement IMW '01**

Publisher: ACM Press

Full text available: pdf(2.51 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Content distribution networks (CDNs) are a mechanism to deliver content to end users on behalf of origin Web sites. Content distribution offloads work from origin servers by serving some or all of the contents of Web pages. We found an order of magnitude increase in the number and percentage of popular origin sites using CDNs between November 1999 and December 2000. In this paper we discuss how CDNs are commonly used on the Web and define a methodology to study how well they perform. A performance ...

16 Provisioning: Efficient and robust streaming provisioning in VPNs



Z. Morley Mao, David Johnson, Oliver Spatscheck, Jacobus E. van der Merwe, Jia Wang

May 2003 **Proceedings of the 12th international conference on World Wide Web WWW '03**

Publisher: ACM Press

Full text available: pdf(1.06 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Today, most large companies maintain virtual private networks (VPNs) to connect their remote locations into a single secure network. VPNs can be quite large covering more than 1000 locations and in most cases use standard Internet protocols and services. Such

VPNs are implemented using a diverse set of technologies such as Frame Relay, MPLS, or IPSEC to achieve the goal of privacy and performance isolation from the public Internet. Using VPNs to distribute live content has recently received treme ...

Keywords: VPNs, streaming server placement

17 Session 8: systems support for multimedia: Cost-effective streaming server implementation using Hi-tactix



Damien Le Moal, Tadashi Takeuchi, Tadaaki Bandoh

December 2002 **Proceedings of the tenth ACM international conference on Multimedia MULTIMEDIA '02**

Publisher: ACM Press

Full text available: pdf(271.85 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

High performance and high quality for continuous media stream delivery needed by streaming server systems cannot be achieved efficiently using general-purpose operating systems, due to the overhead of the I/O mechanism implementation generally used. Special OS combined with powerful hardware can deliver better performance and quality but increases development complexity and deployment costs. The External I/O Engine Architecture adopts a hybrid approach, implementing streaming engines using the s ...

Keywords: audio/video streaming, operating system, quicktime, real-time

18 Special issue on the PAPA 2002 workshop: On the stability of network distance estimation



Yan Chen, Khian Hao Lim, Randy H. Katz, Chris Overton

September 2002 **ACM SIGMETRICS Performance Evaluation Review**, Volume 30 Issue 2

Publisher: ACM Press

Full text available: pdf(947.01 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Overlay network distance monitoring and estimation system can benefit many new applications and services, such as peer-to-peer overlay routing and location. However, there is a lack of such scalable system with small overhead, good usability, and good distance estimation accuracy and stability. Thus we propose a scalable overlay distance monitoring system, *Internet Iso-bar*, which clusters hosts based on the similarity of their perceived network distance, with no assumption about the under ...

19 Resource management with hoses: point-to-cloud services for virtual private networks



N. G. Duffield, Pawan Goyal, Albert Greenberg, Partho Mishra, K. K. Ramakrishnan, Jacobus E. van der Merwe

October 2002 **IEEE/ACM Transactions on Networking (TON)**, Volume 10 Issue 5

Publisher: IEEE Press

Full text available: pdf(425.44 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

As IP technologies providing both tremendous capacity and the ability to establish dynamic security associations between endpoints emerge, virtual private networks (VPNs) are going through dramatic growth. The number of endpoints per VPN is growing and the communication pattern between endpoints is becoming increasingly hard to predict. Consequently, users are demanding dependable, dynamic connectivity between endpoints, with the network expected to accommodate any traffic matrix, as long as the ...

Keywords: point-to-cloud, point-to-multipoint, quality of service, service level agreements

20 Applications and architecture: P2Cast: peer-to-peer patching scheme for VoD service ☐



Yang Guo, Kyoungwon Suh, Jim Kurose, Don Towsley

May 2003 **Proceedings of the 12th international conference on World Wide Web WWW '03**

Publisher: ACM Press

Full text available: pdf(253.27 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Providing video on demand (VoD) service over the Internet in a scalable way is a challenging problem. In this paper, we propose P2Cast - an architecture that uses a peer-to-peer approach to cooperatively stream video using patching techniques, while only relying on unicast connections among peers. We address the following two key technical issues in P2Cast: (1) constructing an application overlay appropriate for streaming; and (2) providing continuous stream playback (without glitches) in the fa ...

Keywords: patching, peer-to-peer networks, performance evaluation, video on-demand service

Results 1 - 20 of 38

Result page: [1](#) [2](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)